arXiv:1608.03226v2 [math.CO] 17 Aug 2016

Drift Analysis and Evolutionary Algorithms Revisited

Johannes Lengler, Angelika Steger
ETH Zürich, Department of Computer Science, Switzerland

In memory of Jiří Matoušek

Abstract

One of the easiest randomized greedy optimization algorithms is the following evolutionary algorithm which aims at maximizing a boolean function $f : \{0,1\}^n \to \mathbb{R}$. The algorithm starts with a random search point $\xi \in \{0,1\}^n$, and in each round it flips each bit of ξ with probability c/n independently at random, where $c > 0$ is a fixed constant. The thus created offspring ξ' replaces ξ if and only if $f(\xi') > f(\xi)$. The analysis of the runtime of this simple algorithm on monotone and on linear functions turned out to be highly non-trivial. In this paper we review known results and provide new and self-contained proofs of partly stronger results.

1 Introduction

Evolutionary algorithms have a long and successful tradition in solving real world optimization problems. Despite much effort the theoretical understanding of these algorithms, however, is still very limited, even for seemingly 'trivial' versions. One of the easiest evolutionary algorithms aims at maximizing a pseudo-boolean function $f : \{0,1\}^n \to \mathbb{R}$. This algorithm starts with a random search point $\xi \in \{0,1\}^n$, and in each round it flips each bit of ξ with probability (or *mutation rate*) c/n independently, where $c > 0$ is the so-called *mutation parameter*. The thus created *offspring* ξ' replaces ξ if and only if $f(\xi') > f(\xi)$. Naturally, the interesting question is to determine the *optimization time* of f, i.e., in the number of rounds until a maximum of f is found (in expectation). Even for the seemingly trivial case that f is a strictly monotone function (that is, $f(\xi) > f(\xi')$ for all ξ, ξ' so that $\xi \neq \xi'$ and $\xi_i \geq \xi'_i$ for all $1 \leq i \leq n$) determining the asymptotic optimization time turned out to be far from trivial. One of the first rigorous results in this direction was [15] who determined the optimization time for the case that $f(\xi) = \sum_{i=1}^{n} \xi_i$. For general linear functions it required substantial efforts [4–6,9,10] until Doerr and Goldberg [2] finally showed that the optimization time is $\Theta(n \log n)$ for all constant mutation parameters $c > 0$. For general monotone functions it is easy to see, cf. e.g. [3], that the optimization time is $\Theta(n \log n)$ for all constants $0 < c < 1$. However, as it turned out, this is not necessarily so for larger mutation parameters. Doerr et al. [3] showed that there are monotone functions such that for $c > 16$ the algorithm takes exponential time.

In this paper we provide short and elegant proofs for various drift theorems that will allow us to give simple proofs for the above statements which are partly stronger than previous results. The exact results can be found in Sections 4.1 and 4.2. The paper is completely self-contained.

2 Basic Tools

The main tool in the proofs (both in previous proofs and in our's) is drift analysis. Going back to the seminal paper of Hajek [7], drift analysis has been developed as a convenient tool to estimate the expected hitting time of algorithms and random processes (see [13] for an overview). In this section we collect various known statements and provide short and self-contained proofs. Our emphasis is on elegance of proofs rather than obtaining general statements. We always assume the following setup. We study a Markov chain $(X_t)_{t \in \mathbb{N}_0}$ over some state space $\mathcal{S}$, which we sometimes restrict to subsets of $\{0\} \cup [1, \infty)$, or simply to $\mathcal{S} = \mathbb{N}$. With T we denote the random variable that denotes the earliest point in time t such that $X_t = 0$. We are usually interested in determining the expectation $\mathbb{E}[T]$ as a function of the

initial state $X_0 = n$. As we are only interested in the first time where X_t reaches zero we may (and do) assume without loss of generality that $X_{T+1} = X_{T+2} = \ldots = 0$.

Our first theorem is a reformulation of Wald's equation, cf. also He and Yao [8].

Theorem 1. *Let $(X_t)_{t\in\mathbb{N}_0}$ be a Markov chain with state space $\mathcal{S} \subseteq [0,\infty)$ and assume $X_0 = n$. If there exists $c > 0$ such that for all $x \in \mathcal{S}, x > 0$ we have*

$$\mathbb{E}[X_{t+1} \mid X_t = x] \leq x - c, \tag{1}$$

then

$$\mathbb{E}[T] \leq \frac{n}{c}.$$

Proof. One easily checks that the condition of theorem implies that for all $t \in \mathbb{N}_0$

$$\mathbb{E}[X_{t+1} \mid T > t] \leq \mathbb{E}[X_t \mid T > t] - c. \tag{2}$$

By conditioning on whether $T > t$ to $T \leq t$ we obtain

$$\mathbb{E}[X_t] = \Pr[T > t] \cdot \mathbb{E}[X_t \mid T > t] + \Pr[T \leq t] \cdot 0 \;=\; \Pr[T > t] \cdot \mathbb{E}[X_t \mid T > t]. \tag{3}$$

Proceeding similarly for X_{t+1} we obtain

$$\begin{aligned} \mathbb{E}[X_{t+1}] \;&=\; \Pr[T > t] \cdot \mathbb{E}[X_{t+1} \mid T > t] + \Pr[T \leq t] \cdot 0 \\ &\overset{(2)}{\leq} \Pr[T > t] \cdot (\mathbb{E}[X_t \mid T > t] - c) \\ &\overset{(3)}{=} \mathbb{E}[X_t] - \Pr[T > t] \cdot c. \end{aligned} \tag{4}$$

Since T is a random variable that takes values in $\mathbb{N}$, we may write $\mathbb{E}[T] = \sum_{t=0}^{\infty} \Pr[T > t]$. For simplicity assume that $\lim_{t\to\infty} \mathbb{E}[X_t]$ exists. Then

$$c \cdot \mathbb{E}[T] = \sum_{t=0}^{\infty} c \Pr[T > t] \overset{(4)}{\leq} \sum_{t=0}^{\infty} (\mathbb{E}[X_t] - \mathbb{E}[X_{t+1}]) \;=\; \mathbb{E}[X_0] - \underbrace{\lim_{t\to\infty} \mathbb{E}[X_t]}_{\geq 0} \;\leq\; n, \tag{5}$$

which proves the claim of the theorem. If $\lim_{t\to\infty} \mathbb{E}[X_t]$ does not exist then (5) still follows with $\liminf$ instead of lim.

As a side remark, note that if Condition (2) is replaced by the reverse condition $\mathbb{E}[X_{t+1} \mid T > t] \geq \mathbb{E}[X_t \mid T > t] - c$, then the same proof shows that $\mathbb{E}[T] \geq \frac{1}{c}(n - \liminf_{t\to\infty} \mathbb{E}[X_t])$.* □

Note that the assumption of Theorem 1 cannot hold for Markov chains whose state space $\mathcal{S}$ contains values arbitrarily close to zero. The following theorem allows to resolve this problem, as it shows that we can partition the analysis of such Markov chains into two phases: one coming 'close' to zero, the other actually reaching zero.

Theorem 2. *Let $(X_t)_{t\in\mathbb{N}_0}$ be a Markov chain with state space $\mathcal{S} \subseteq [0,\infty)$ and assume $X_0 = n$. Furthermore let $C > 0$ be some positive constant and denote by T_C the earliest point in time t such that $X_t \leq C$. Assume that there exist constants $p_0 > 0$ and $B > 0$ such that for all $c \in \mathcal{S}$ with $0 < c < C$ we have*

$$\Pr[X_{t+1} = 0 \mid X_t = c] \geq p_0$$

and

$$\sum_{x\in\mathcal{S}, x>C} \Pr[X_{t+1} = x \mid X_t = c] \cdot \mathbb{E}[T_C \mid X_0 = x] \leq B.$$

Then $\mathbb{E}[T] = \mathbb{E}[T_C] + O(1)$.

*So either $\mathbb{E}[T] = \infty$, or $\mathbb{E}[T]$ is a finite value which satisfies the inequality.

Proof. Let $\mathcal{S}_C := \{s \in \mathcal{S} \mid s \leq C\}$. By the definition of T_C the expectation of the number of steps until we reach a state in $c \in \mathcal{S}_C$ is $\mathbb{E}[T_C]$. So assume now that we are in some nonzero state $c \in \mathcal{S}_C$. With probability at least p_0 we jump right into zero in the next step. Otherwise we either end up in some (possibly different) state in $\mathcal{S}_C$ or in some state $x > C$. In the latter case we need, due to our assumption, in expectation at most B steps to return to a state $\mathcal{S}_C$. We conclude that we have in expectation at least every $(B+1)$st step a chance of jumping to zero with probability at least p_0. As the expectation of a geometric random variable with success probability p_0 is $1/p_0$ this implies that, once we reach a state within $\mathcal{S}_C$ for the first time, the expectation for the number of additional steps to reach zero is bounded from above by $(B+1)/p_0$, as claimed. □

Consider a Markov chain on the non-negative integers which is defined as follows. From all states $x > 0$ we move to one of the states $0, 1, \ldots, x$ uniformly at random. How long does it take till we reach zero? Intuitively, we expect to half the distance towards zero in each step, so we expect that it will take $\log_2 n$ steps, if we start in n. What if we move uniformly to a state in $0, 1, \ldots, ax$? For which $a > 1$, if at all, do we have time $O(\log n)$ till we reach zero? We will give the (surprising) answer for this in the next section. The following theorem will turn out to be very useful. The clue is to consider a Markov chain in which the distance to zero is rescaled.

Theorem 3. *Let $(X_t)_{t\in\mathbb{N}_0}$ be a Markov chain with state space $\mathcal{S} \subseteq \mathbb{N}$ and with $X_0 = n$. Let $C \in \mathbb{N}$ be some positive constant and denote by T_C the earliest point in time t such that $X_t \leq C$. Assume furthermore that there exists a constant $c > 0$ and a function $g : \mathbb{N}_0 \to \mathbb{R}$ such that $g(x) = 0$ for all $x \leq C$ and $g(x) > 0$ for all $x > C$ and*

$$\sum_{n>C} \Pr[X_{t+1} = n \mid X_t = x] \cdot g(n) \leq g(x) - c \qquad \text{for all } x \in \mathbb{N}, x > C.$$

Then $\mathbb{E}[T_C] \leq g(n)/c$.

Proof. Let $Y_t := g(X_t)$, if $X_t > C$ and $Y_t = 0$ otherwise. Then the assumption of the theorem implies that for all $x > C$:

$$\mathbb{E}[Y_{t+1} \mid Y_t = g(x)] = \sum_{n>C} \Pr[X_{t+1} = n \mid X_t = x] \cdot g(n) \leq g(x) - c.$$

We can thus apply Theorem 1 to the Markov chain (Y_t) to conclude that $\mathbb{E}[\widetilde{T}] \leq g(n)/c$, where $\widetilde{T}$ is the time until the Markov chain (Y_t) reaches zero. The observation that, by construction of the chain (Y_t), we have $\widetilde{T} \equiv T_C$ concludes the proof. □

The next theorem is a simplified version of the previous one that is often easier to apply. The present formulation is from Mitavskiy, Rowe, and Cannings [14], see also [12, Theorem 4.6].

Theorem 4. *Let $(X_t)_{t\in\mathbb{N}_0}$ be a Markov chain with state space $\mathcal{S} \subseteq \{0\} \cup [1, \infty)$ and with $X_0 = n$. Let T be the random variable that denotes the earliest point in time $t \geq 0$ such that $X_t = 0$. Suppose furthermore that there is a positive, increasing function $h : \mathbb{N} \to \mathbb{R}_{>0}$ such that for all $x \in \mathcal{S}$, $x > 0$ we have*

$$\mathbb{E}[X_{t+1} \mid X_t = x] \leq x - h(x).$$

Then

$$\mathbb{E}[T] \leq \frac{1}{h(1)} + \int_1^n \frac{1}{h(u)} du.$$

Proof. Let

$$g(x) := \frac{1}{h(1)} + \int_1^x \frac{1}{h(u)} du \quad \text{for } x \in \mathcal{S}, x > 0 \qquad \text{and } g(0) := 0.$$

We claim that $g(x) - g(y) \geq \frac{x-y}{h(x)}$ for all $x, y \in N$ such that $x > 0$. To see this, observe that

$$g(x) - g(y) = \begin{cases} \int_y^x \frac{1}{h(u)} du \geq (x-y) \cdot \frac{1}{h(x)} & \text{if } x > y > 0, \\ \frac{1}{h(1)} + \int_1^x \frac{1}{h(u)} du \geq \frac{x}{h(x)} & \text{if } x > y = 0, \\ 0 & \text{if } x = y > 0, \\ -\int_x^y \frac{1}{h(u)} du \geq -(y-x) \cdot \frac{1}{h(x)} & \text{if } 0 < x < y. \end{cases}$$

(The inequalities follow from the fact that we assumed that h is increasing and positive.)

Now let us consider the random variable $Y_t := g(X_t)$. The above inequality for g implies that for all $x \in \mathbb{N}$ we have

$$\mathbb{E}[g(X_t) - g(X_{t+1}) \mid X_t = x] \geq \mathbb{E}\left[\frac{X_t - X_{t+1}}{h(X_t)} \mid X_t = x\right] \geq 1,$$

where the last inequality is just a reformulation of the assumption of the theorem.

Hence, we have $\mathbb{E}[Y_{t+1} \mid Y_t = g(x)] \leq g(x) - 1$, and the claim of the theorem thus follows by applying Theorem 1 to $Y_t := g(X_t)$. □

An special case of Theorem 4 is the case $h(x) = \delta x$; we then get $\mathbb{E}[T] \leq \delta^{-1}(1 + \log n)$. In fact, a straightforward application of Markov's inequality even supplies us with exponential tail bounds, as was first noticed in [4].

Theorem 5. *Let $(X_t)_{t\in\mathbb{N}_0}$ be a Markov chain with state space $\mathcal{S} \subseteq \{0\} \cup [1, \infty)$ and with $X_0 = n$. Assume that there is $\delta > 0$ such that for all $x \in \mathcal{S}$,*

$$\mathbb{E}[X_{t+1} \mid X_t = x] \leq (1 - \delta)x.$$

Then

$$\Pr\left[T > \left\lceil \frac{\log n + k}{|\log(1 - \delta)|} \right\rceil\right] \leq e^{-k}.$$

Proof. Let $t := \lceil (\log n + k)/|\log(1 - \delta)| \rceil$. Since $\mathbb{E}[X_t] \leq (1 - \delta)^t X_0 \leq e^{-k}$, the claim follows from Markov's inequality, $\Pr[T > t] = \Pr[X_t \geq 1] \leq e^{-k}$. Note that here we used $\mathcal{S} \subseteq \{0\} \cup [1, \infty)$; replacing the (somewhat arbitrary) boundary 1 by some other constant would give a multiplicative constant in front of the probability e^{-k}. □

An important implication of Theorem 5 is that if the probability of making large jumps is geometrically bounded, then also an additive drift suffices to get exponential tail bounds. Moreover, it is exponentially unlikely to cross an area with *negative* drift in less than exponentially many steps. This idea again goes back to Hajek [7] and was first proven in this simplified form in [16]. However, that proof is still based on some results from Hajek's paper. Here we give a short and fully self-contained proof.

Theorem 6. *Let $(X_t)_{t\in\mathbb{N}_0}$ be a Markov chain with state space $\mathcal{S} \subseteq \mathbb{R}$. Suppose there exist constants $a, b, \delta, \varepsilon, r > 0$ such that for all $t \geq 0$ the following conditions hold:*

1. $\mathbb{E}[X_{t+1} - X_t \mid an < X_t] \leq -\varepsilon$,
2. $\Pr[|X_{t+1} - X_t| \geq j] \leq r(1 + \delta)^{-j}$ *for all* $j \in \mathbb{N}$.

Let $T_a := \min\{t \geq 0 : X_t \leq an \mid X_0 = bn\}$ and $T_b := \min\{t \geq 0 : X_t \geq bn \mid X_0 = an\}$. Then for any constant $c_1 > 0$,

(a) $\Pr[T_a \geq \frac{(1+c_1)(b-a)n}{\varepsilon}] \leq e^{-\Omega(n)}$.

Moreover, there is a constant $c_2 > 0$ such that

(b) $\Pr[T_b \leq e^{c_2 n}] = e^{-\Omega(n)}$.

Proof. We may assume $c_1 < 1/2$. Let $0 < \eta < 1$ be so small that $e^\eta < (1 + \delta)^{1/2}$ holds. We claim that if η is sufficiently small, then we can choose $j_0 > 0$ such that $e^{\eta j} = 1 + \eta j \pm c_1\eta\varepsilon/6$ is satisfied for all $-j_0 \leq j \leq j_0$ (where we use $x = y \pm \varepsilon$ as shorthand for $x \in [y - \varepsilon, y + \varepsilon]$), and such that at the same time

$$\sum_{j=j_0+1}^{\infty} r(1 + \delta)^{-j/2} < c_1\eta\varepsilon/24 \qquad \text{and} \qquad \sum_{j=j_0+1}^{\infty} jr(1 + \delta)^{-j} < c_1\eta\varepsilon/24 \tag{6}$$

holds. Indeed, by setting (rather arbitrarily) $j_0 = \eta^{-1/3}$ we satisfy all three conditions if $\eta = \eta(\delta, \varepsilon, c_1, r)$ is small enough.

Consider $Y_t := e^{\eta X_t}$. For any $x > an$ we have

$$\begin{aligned}
\mathbb{E}[Y_{t+1} - Y_t \mid X_t = x] &= e^{\eta x}\mathbb{E}[e^{\eta(X_{t+1}-x)} - 1 \mid X_t = x] \\
&= e^{\eta x}\Bigg(\sum_{j=-j_0}^{j_0} \Pr[X_{t+1} - X_t = j \mid X_t = x] \underbrace{(e^{\eta j} - 1)}_{=\eta j \pm c_1\eta\varepsilon/6} \pm 2 \cdot \underbrace{\sum_{j=j_0+1}^{\infty} r(1+\delta)^{-j}(e^{\eta j}+1)}_{\leq c_1\eta\varepsilon/12 \text{ by (6)}}\Bigg) \\
&= e^{\eta x}\Bigg(\sum_{j=-j_0}^{j_0} \Pr[X_{t+1} - X_t = j \mid X_t = x]\eta j \pm c_1\eta\varepsilon/3\Bigg) \\
&\overset{(6),\text{Cond.2}}{=} e^{\eta x}(\mathbb{E}[X_{t+1} - X_t \mid X_t = x]\eta \pm c_1\eta\varepsilon/2) < -e^{\eta x}\underbrace{\eta\varepsilon(1 - c_1/2)}_{=:\gamma}.
\end{aligned}$$

Recalling that $Y_t = e^{\eta x}$ for $X_t = x$, we thus find that Y_t has a multiplicative drift with factor $1 - \gamma < 1 - \eta\varepsilon/(1+c_1)$. The first statement now follows immediately from Theorem 5. For the second statement, consider the event $\mathcal{E}$ that for the first $e^{c_2 n}$ points in time where $X_t \leq an$ we have $X_{t+1} \leq X_t + (b-a)n/2$. Note that $\Pr[\neg\mathcal{E}] \leq e^{c_2 n} \cdot r(1-\delta)^{n(b-a)/2} = e^{-\Omega(n)}$, if we choose $c_2 > 0$ sufficiently small. From now on we may thus condition on $\mathcal{E}$. Consider some time t_1 at which we cross the lower bound from below, i.e., $X_{t_1-1} \leq an$ and $X_{t_1} > an$. Let t_2 be the next point in time where we drop back below an, so $t_2 = \min\{t > t_1 \mid X_{t_2} \leq an\}$. Note that the event $\mathcal{E}$ implies that for all $t_1 < e^{c_2 n}$ we have $X_{t_1} < n(a+b)/2$. We now show that it is very unlikely that we reach b between t_1 and $\max\{t_2, e^{c_2 n}\}$. In fact, this can only happen if $Y_t \geq e^{\eta b}$. Thus, using the multiplicative drift of Y_t and Markov's inequality, for all $t_1 < t < \max\{t_2, e^{c_2 n}\}$ and $an < y \leq (a+b)n/2$,

$$\Pr[X_t > bn \mid X_{t_1} = y] \leq \frac{\mathbb{E}[Y_t \mid Y_{t_1} = e^{\eta y}]}{e^{\eta bn}} \leq \frac{(1-\gamma)^{t-t_1}e^{\eta y}}{e^{\eta bn}} \leq (1-\gamma)^{t-t_1}e^{-\eta(b-a)n/2}.$$

Union bound over all $t_1 < t < \max\{t_2, e^{c_2 n}\}$ shows that (conditioned on $\mathcal{E}$) the probability that bn is reached in the interval $[t_1, t_2]$ is $e^{-\Omega(n)}$. A second union bound over all $t_1 < e^{c_2 n}$ thus shows the second claim for c_2 sufficiently small. □

3 Example: Random Decline

With Theorem 4 at hand we can now analyse the Markov chain that, from state x, moves to a state uniformly at random within $\{0, \ldots, \lfloor ax \rfloor\}$. Clearly, $\mathbb{E}[X_{t+1} \mid X_t = x] \leq \frac{a}{2}x$ and we can apply Theorem 4 to deduce that the number of steps T until we reach zero is $O(\log n)$ for all $a < 2$. For $a > 2$ it may seem that $\mathbb{E}[T] = \infty$, as we have a drift to the right. Perhaps surprisingly, this is not true. In fact, we have the following theorem.

Theorem 7. *For the Markov chain defined above with (constant) parameter a, we have $\mathbb{E}[T] = O(\log n)$ if and only if $a < e$, where $e = 2.718..$ is the Euler constant.*

Proof. Assume $a < e$. Let $g(x) = \log(x)$ for $x > C$ (where C is some constant depending on a that we will fix below) and $g(x) = 0$ for all $x \leq C$. Then

$$\sum_{m>C} \Pr[X_{t+1} = m \mid X_t = x] \cdot g(m) = \frac{1}{\lfloor ax \rfloor + 1} \sum_{m=C+1}^{\lfloor ax \rfloor} \log(m) \leq \log(x) - 1 + \log(a) + \frac{1}{aC+1}.$$

As $\log(a) < 1$ for all $a < e$ we can thus find for all $a < e$ a constant C so that we can apply Theorem 3 for some constant $c > 0$ to deduce that $\mathbb{E}[T_C] \leq O(\log n)$. Theorem 2 then implies that we also have $\mathbb{E}[T] = O(\log n)$. Indeed, the assumptions of Theorem 2 are easily seen to be fulfilled for $p_0 = 1/(C+1)$ and $B = \max_{C+1 \leq n \leq aC} \mathbb{E}[T_C \mid X_0 = n]$ (which is finite, as we maximize over a finite number of expectations and we have already shown that these expectations are finite).

For the other direction, we restrict ourselves for concreteness to $a = e$. (The case $a > e$ then follows easily by coupling it to the case $a = e$.) For the sake of contradiction we assume that $\mathbb{E}[T] = O(\log n)$,

so assume that there is $D > 0$ such that $\mathbb{E}[T] \leq D \log n$ for all n. This implies in particular that $\mathbb{E}[T]$ is finite. Note that we may assume $D \geq 2$, and that D is an integer. Let $g(x) = 2D \log(x/D^2)$ for $x > D^2$, $g(x) = 1$ for $1 \leq x \leq D^2$, and $g(x) = 0$ for $x = 0$. We set $Y_t := g(X_t)$, where X_t is the state in the t-th step. Note that $X_t = 0$ if and only if $Y_t = 0$. We claim that the reverse of Condition (2) in Theorem 1 is satisfied for Y_t for $c = 1$. To see this consider some $y > 0$ of the form $g(x)$, where $x \in \mathbb{N}^+$. We want to show that $\mathbb{E}[Y_{t+1} \mid Y_t = y] \geq y - 1$. For $x \leq D^2$ this is trivial since then $y = g(x) \leq 1$. So assume that $x > D^2$. We compute similarly as above (using that $s \cdot \log(s/(eD^2)$ is an antiderivative of $\log(s/D^2)$ and that $\lfloor ex \rfloor \geq ex - 1$ implies $\log(\lfloor ex \rfloor/(eD^2)) \geq \log(x/D^2) - 2/(ex)$) that

$$\begin{aligned}
\mathbb{E}[Y_{t+1} \mid Y_t = y] = \mathbb{E}[g(X_{t+1}) \mid X_t = x] &\geq \frac{1}{\lfloor ex \rfloor + 1} \sum_{m=D^2}^{\lfloor ex \rfloor} (2D \log(m/D^2)) \\
&= \frac{\lfloor ex \rfloor}{\lfloor ex \rfloor + 1} \left(2D \log \left(\frac{\lfloor ex \rfloor}{eD^2} \right) + \frac{2D^3}{\lfloor ex \rfloor} \right) \\
&\geq 2D \log(x/D^2) - 1 \;=\; y - 1.
\end{aligned}$$

Since this inequality is true for all $y = g(x) > 0$, $x \in \mathbb{N}^+$, we may use the reverse version of Theorem 1, cf. the remark at the end of its proof. Hence, for $n > D^2$,

$$\mathbb{E}[T] \geq Y_0 - \liminf_{t \to \infty} \mathbb{E}[Y_t] = 2D \log(n/D^2) - \liminf_{t \to \infty} \mathbb{E}[Y_t].$$

Since we assumed $\mathbb{E}[T] \leq D \log n$, the last inequality implies for all sufficiently large n and t

$$\mathbb{E}[Y_t] \geq D \log(n)/2.$$

Let $p_t := \Pr[Y_t > 0] = \Pr[T > t]$, and let $E_t := \mathbb{E}[Y_t \mid Y_t > 0]$. Then $p_t E_t = \mathbb{E}[Y_t] \geq D \log(n)/2$ for sufficiently large n and t. On the other hand, since X_t can increase in each step by at most a factor of e, we can bound $X_t \leq ne^t$, and thus $E_t \leq 2D \log(ne^t/D^2)$. Combining both bounds, we obtain

$$p_t \geq \frac{D \log n}{2E_t} \geq \frac{D \log n}{4D \log(ne^t/D^2)} = \frac{D \log n}{4Dt + 4D \log(n/D^2)}.$$

However, this implies that

$$\mathbb{E}[T] = \sum_{t \geq 0} p_t = \infty,$$

contradicting our assumption that $\mathbb{E}[T]$ is finite. This proves that $\mathbb{E}[T] = \omega(\log n)$ for $a = e$. □

4 Evolutionary Algorithm

In this section we (re)prove some fundamental properties of the following classical algorithm [1] for optimizing a *fitness function* $f : \{0,1\}^n \to \mathbb{R}$, which we call *evolutionary algorithm* or *EA*[†]. The algorithm starts with a search point $\xi^0 \in \{0,1\}^n$ that we usually assume to be chosen uniformly at random. It then proceeds in rounds. In each round it flips each bit of the current search point ξ with probability (or *mutation rate*) c/n, where $c > 0$ is the *mutation parameter*. The thus created *offspring* ξ' replaces ξ if and only if $f(\xi') > f(\xi)$. We are interested in the *optimization time* of f, i.e., in the number of rounds until a (global) maximum of f is found. Note that we assume that the algorithm has access to the function f only via an oracle. That is, the algorithm can query the function value $f(\xi)$ for a given ξ, but has no further knowledge about the function. In particular, the algorithm thus has no chance to 'know' which x maximizes the function before having queried all 2^n values. This is the reason for the above definition of the optimization time: we count the number of rounds until the algorithm first queries a global maximum. From then on the algorithm may continue asking queries, but clearly, the function value will never change again.

Throughout this section, we will use the following notation. For all $t \geq 0$, we denote by $\xi^t \in \{0,1\}^n$ the search point after t rounds of the algorithm. For any non-empty set of indices $I \subseteq [n]$, let $d(I,t) :=$

[†]In the EA community it is called (1+1)-Evolutionary Algorithm.

$|\{i \in I \mid \xi_i^t = 0\}|/|I|$ be the density of zero bits in the I-substring of ξ^t. We write $d(I)$ instead of $d(I,t)$ if t is clear from the context.

By the *no free lunch* principle, for arbitrary functions f there is no hope that the algorithm will be fast. For example, for the function $f(\xi)$ that is one for $\xi = x_0$ and zero otherwise, *any* algorithm will have exponential optimization time whp[‡]. Nevertheless algorithm EA gives good results in many practical applications, indicating the algorithm should be faster for a suitably restricted class of functions. In this section, we will study two classes of functions: *strictly monotone functions*, that is, functions f that satisfy $f(\xi) > f(\xi')$ for all ξ, ξ' so that $\xi \neq \xi'$ and $\xi_i \geq \xi'_i$ for all $1 \leq i \leq n$; and *linear functions*, i.e., functions of the form $f(\xi) = \sum_{i=1}^n a_i \xi_i$ with weights $a_1, \ldots, a_n \in \mathbb{R}$. To avoid trivialities, we will always assume that all a_i are non-zero. Moreover, by symmetry of the EA we may even assume that the weights are sorted and positive, $a_1 \geq \ldots \geq a_n > 0$. Note that with these assumptions every linear function is in particular strictly monotone.

For strictly monotone functions it easily follows from the drift theorems in Section 2 that the optimization time of algorithm EA is $\Theta(n \log n)$ for all mutation parameters $c < 1$, see also [3].

Theorem 8. *For every constant $c < 1$ and all strictly monotone functions f, the optimization time of the algorithm EA with mutation rate c/n is $\Theta(n \log n)$ in expectation and whp.*

Proof. For the lower bound, simply observe that whp the initial string has at least $n/3$ zero bits, and that each of these bits needs to be flipped at least once to reach the optimum. Therefore, we need to flip at least $\Omega(n \log n)$ bits in expectation and whp (since this is a coupon collector process), and thus we also need $\Omega(n \log n)$ rounds to achieve so many flips.

For the upper bound, let $X_t := |\{i \in [n] : \xi_i^t = 0\}|$ be the number of zero bits in the t-th step of the algorithm. We want to bound $\mathbb{E}[X_{t+1} \mid X_t = x]$. Assume we flip r zero bits to a one and s one bits to a zero. If $r = 0$, then we reject the offspring ξ' and thus $X_{t+1} = X_t$. If $r > 0$ and $s = 0$ then we accept ξ' and thus $X_{t+1} \leq X_t - 1$, with room to spare. For the remaining case $r > 0$ and $s \geq 1$ we use the bound $X_{t+1} \leq X_t + s - 1$, which is true regardless whether we accept or reject ξ'.

Denote by $p_{r,0}$ the probability that $r = 0$. Then the above observations imply

$$\mathbb{E}[X_{t+1} \mid X_t = x] \leq p_{r,0} x + (1 - p_{r,0})(x - 1 + \mathbb{E}[s]) = x - (1 - p_{r,0})(1 - \mathbb{E}[s]).$$

Note that $p_{r,0}$ is the probability that a binomially distributed random variable with parameters x and $p = c/n$ is equal to zero. Here the following inequality about the binomial distribution is helpful:

$$\Pr[\mathrm{Bin}(n,p) > 0] \geq \frac{np}{1+np} \qquad \text{for all } n \in \mathbb{N} \text{ and } 0 \leq p \leq 1. \tag{7}$$

(Easy to check by induction on n.) With (7) and the observation that $\mathbb{E}[s] = (n - x) \cdot (c/n) \leq c$ we deduce

$$\mathbb{E}[X_{t+1} \mid X_t = x] \leq x - \frac{xc/n}{1 + xc/n}(1 - c) \leq x - \frac{xc(1-c)}{n(1+c)}$$

and we can apply Theorem 4 with $h(x) = xc(1-c)/(n(1+c))$ in order to deduce that $\mathbb{E}[T] = O(n \log n)$. Moreover, by Theorem 5 (with $\delta = c(1-c)/(n(1+c))$) the same bound holds whp. □

We will see in the next section that the restriction $c < 1$ is essential: for $c > 2.2$ the running time is not $O(n \log n)$, but actually exponential in n. Before we proceed to that proof we state a lemma on linear functions that we will turn out to be very useful in Section 4.1 as well as Section 4.2.

Lemma 9. *Consider the algorithm EA on a linear function, $f(x) = \sum_{i=1}^n a_i x_i$, and assume that $a_i \geq a_j$ for $i \leq j$. Let $\alpha, \beta, \varepsilon > 0$ be any constants and let $I, J \subseteq [n]$ be two sets of indices such that $\max I < \min J$ and such that $|I| = \alpha n$ and $|J| = \beta n$. Then there exists $C > 0$ such that for any initial value ξ^0 of the algorithm, whp we have $d(I,t) \leq d(J,t) + \varepsilon$ for all $t \in [Cn, e^{n/C}]$. The interval may be replaced by $[0, e^{n/C}]$ if $d(I,0) \leq d(J,0) + \varepsilon/2$.*

[‡] *with high probability*, i.e., with probability tending to one as $n \to \infty$

Proof. We first consider a single step, so fix $t \geq 0$ and ξ^t. Let $i < j$ be two positions such that $\xi_i^t = \xi_j^t = 0$. We claim that $\Pr[\xi_i^{t+1} = 1] \geq \Pr[\xi_j^{t+1} = 1]$. To see this assume first that all bit flips except the ith and the jth are fixed. Then the following holds. If the bit flips for the ith and the jth bit are identical then we change either neither of the two bits or both. Thus, the only interesting case is if exactly one of these two bits is flipped. Here we observe the following: if the case that the jth bit is flipped is accepted then so is the case that the ith bit is flipped (due to the fact that we assumed $a_1 \geq a_2 \geq \ldots \geq a_n$). A few moments of thought show that this already proves our claim.

Similarly we deduce that if $k < \ell$ are two positions such that $\xi_k^t = \xi_\ell^t = 1$, then we have $\Pr[\xi_k^{t+1} = 0] \leq \Pr[\xi_\ell^{t+1} = 0]$.

Now we can prove the lemma. By Theorem 6 it suffices to show that $n \cdot (d(I) - d(J))$ has a negative drift, whenever $d(I) \geq d(J) + \varepsilon/2$. (The second condition in Theorem 6 is satisfied since I and J have linear size.) So assume that $d(I,t) \geq d(J,t) + \varepsilon/2$ holds for some $t \geq 0$. Let $p_{I0} := \min\{\Pr[\xi_i^{t+1} = 1] : i \in I, \xi_i^t = 0\}$ and $p_{J0} := \max\{\Pr[\xi_j^{t+1} = 1] : j \in J, \xi_j^t = 0\}$ and observe that we know from above that $p_{I0} \geq p_{J0}$. Similarly, we have for $p_{I1} := \max\{\Pr[\xi_k^{t+1} = 0] : k \in I, \xi_k^t = 1\}$ and $p_{J1} := \min\{\Pr[\xi_\ell^{t+1} = 0] : \ell \in J, \xi_\ell^t = 1\}$ that $p_{I1} \leq p_{J1}$. Moreover, observe that $p_{J0} = \Omega(1/n)$, because the probability to flip the ℓ-th bit and no other bit is $\Omega(1/n)$. Thus, the the expected change of $n \cdot (d(I,t) - d(J,t))$ is at most

$$n[(1 - d(I,t))p_{I1} - d(I,t)p_{I0}] - n[[(1 - d(J,t))p_{J1} - d(J,t)p_{J0}] \leq n \underbrace{(d(J,t) - d(I,t))}_{\leq -\varepsilon/2}(p_{I1} + p_{J0}),$$

which is $-\Omega(1)$, as desired. □

4.1 Monotone Functions

We already saw that for $c < 1$ the runtime of algorithm EA is $\Theta(n \log n)$ for all strictly monotone functions. For the natural choice $c = 1$ it is still unknown whether the runtime is $\Theta(n \log n)$. The best known upper bound is $O(n^{3/2})$ due to Jansen [11]. The following theorem explains why this may be not so easy to show, as we do get an exponential optimization time for c's that are not much larger.

Theorem 10. *For every constant $c \geq 2.2$, there is a strictly monotone function such that whp the optimization time of the EA with mutation rate c/n is $e^{\Omega(n)}$.*

A similar theorem was shown by Doerr et al. [3] with $c > 16$ instead of $c > 2.2$.

Proof. The main idea of our proof is the following. Consider n bits and split them in high order bits (first αn, say) and low order bits (the remaining $(1 - \alpha)n$). Now consider the (linear) fitness function $a_i = n$ for $1 \leq i \leq \alpha n$ and $a_i = 1$ otherwise. Then the density of zero bits will decrease faster among the high order bits. In particular, there will be a time when the density among the high order bits is ε, but the density among the low order bits is still $\varepsilon + \delta$, for some $\delta = \delta(\alpha, \varepsilon, c) > 0$. The overall density in the string is then $\varepsilon + (1 - \alpha)\delta > \varepsilon$.

If we now change the fitness function by picking a random set of αn bits as 'new' high order bits, we will start with a density of $\varepsilon + (1 - \alpha)\delta$ among the high order bits and will decrease it during the second round to density ε. Whenever we decrease the number of zero bits among the high order bits, we accept the offspring regardless of the changes of low order bits. In particular, in such rounds the density among the low order bits will *increase* in expectation towards $1/2$. By choosing the constants appropriately, this yields overall a *positive* drift for the density among the low order bits. The drift is actually strong enough that the density will increase from $\varepsilon + (1 - \alpha)\delta$ to $\varepsilon + \delta$ in a very short time – this will finish before the density among high order bits can decrease from $\varepsilon + (1 - \alpha)\delta$ to ε. So after the second round again we are in the situation that the density among high order bits is ε and the density among low order bits is (at least) $\varepsilon + \delta$. We will show that we can play this game for exponentially many rounds.

Formally we proceed as follows. For a given $c \geq 2.2$ we choose $0 < \alpha < 1/2$ which satisfies

$$\alpha c - e^{-(1-\alpha)c} > \frac{\alpha}{1 - \alpha}. \tag{8}$$

A numerical calculation shows that this is possible for all $c \geq 2.2$. We choose a constant $\beta = \beta(\alpha, c)$ which is sufficiently small compared to α (by abuse of notation we denote this by $\alpha \gg \beta$). To provide an overview, the constants in the proof will satisfy

$$1/2 > \alpha \gg \varepsilon \gg \delta \gg \beta \gg \gamma \gg \mu > 0, \tag{9}$$

where γ and μ are constants that will appear later in the proof. The $\gg$-sign hereby indicates the only restriction that we require: the constants are chosen from left to write, so that α satisfies equation (8) and then every constant is suitably small compared to all previously chosen constants.

For $1 \leq i \leq e^{\mu n}$ we choose sets $A_i \subseteq [n]$ of size αn independently and uniformly at random, and we choose sets $B_i \subseteq A_i$ uniformly at random of size βn. We define the *level* $\ell(\xi)$ of a vector $\xi \in \{0,1\}^n$ by

$$\ell(\xi) := \max\{\ell' \in [e^{\mu n}] : |\{j \in B_{\ell'} : \xi_j = 0\}| \leq \varepsilon |B_{\ell'}|\} \quad \text{(with } \ell(\xi) = 0 \text{, if no such } \ell' \text{ exists)}.$$

With this notation at hand we can then define a fitness function $f : \xi \in \{0,1\}^n \to \mathbb{R}$ as follows:

$$f(\xi) = \ell(\xi) \cdot n^{10} + \sum_{i \in A_{\ell(\xi)+1}} \xi_i \cdot n^6 + \sum_{i \notin A_{\ell(\xi)+1}} \xi_i.$$

So the set A_{i+1} describes the set of high order bits at level i, where the level is determined by the sets B_i. One easily checks that this function is monotone. Indeed, assume that ξ is dominated by ξ', i.e., $\xi_i \leq \xi'_i$ for all $1 \leq i \leq n$. Then $\ell(\xi) \leq \ell(\xi')$. If $\ell(\xi) = \ell(\xi')$ then it is obvious that $f(\xi) \leq f(\xi')$ (and we have a strict inequality if $\xi \neq \xi'$). On the other hand, if $\ell(\xi) < \ell(\xi')$ then $f(\xi) \leq \ell(\xi) n^{10} + n^7 < \ell(\xi') n^{10} \leq f(\xi')$, as desired.

Observe also that, by definition of the function f, the algorithm will never accept a change that will decrease the *current level* $\ell(t) := \ell(\xi^t)$. That is, when the current level is i the density of zero bits will stay below ε within B_i (or the level increases).

Note that the definition of ℓ refers to *all* sets B_i, which makes the analysis quite tricky. In order to avoid this we couple the process with a similar process where we only have to choose the sets A_i and B_i as the algorithms progresses. Consider the random variable $\tilde{\ell} = \tilde{\ell}(t)$, which (as we will show later) coincides whp with the current level $\ell = \ell(t)$ for an exponential number of rounds. Initially we set $\tilde{\ell} = 0$, and in each round we increase $\tilde{\ell}$ by one if and only if $|\{j \in B_{\tilde{\ell}+1} : \xi^t_j = 1\}| \geq (1-\varepsilon)|B_{\tilde{\ell}+1}|$. Otherwise we leave $\tilde{\ell}$ unchanged. Clearly, as long as ℓ increases only in steps of one, it coincides with $\tilde{\ell}$. Consider algorithm EA on the time-dependent fitness function

$$\tilde{f}(\xi, t) = \tilde{\ell}(t) \cdot n^{10} + \sum_{i \in A_{\tilde{\ell}(t)+1}} \xi_i \cdot n^6 + \sum_{i \notin A_{\tilde{\ell}(t)+1}} \xi_i.$$

As long as $\tilde{\ell}(t)$ coincides with $\ell(t)$ we have $\tilde{f}(\xi^t, t) = f(\xi^t)$, so the algorithm performs the same steps on both functions as long as $\tilde{\ell}(t) = \ell(t)$. The advantage of $\tilde{f}$ is that we do not have to choose the sets A_i, B_i in advance: we can postpone choosing A_{i+1}, B_{i+1} until we enter the ith level.

Before we now dive into the proof we fix some notations. Let $T := e^{\mu n}$. Our aim is to show that whp the algorithm will run for at least T rounds. With t_i we denote the points in time where $\tilde{\ell}$ increases from $i-1$ to i, i.e., $\tilde{\ell}(t_i - 1) = i - 1$ and $\tilde{\ell}(t_i) = i$. Recall that for a set $S \subseteq [n]$ we denote by $d(S,t)$ the density of zero bits in the set S with respect to the tth search point ξ^t.

Observe that the algorithm EA (with respect to the function $\tilde{f}$) does not need to know A_{i+1} before time t_i. By Chernoff bounds (and a union bound over all $1 \leq t \leq t_i \leq T$) we can thus assume (by making μ sufficiently small) that for an arbitrarily small constant γ we have whp that

$$\text{for all } i \text{ s.t. } t_i \leq T\text{:} \quad d(A_{i+1}, t) = d([n], t) \pm \gamma \quad \text{for all } 1 \leq t \leq t_i. \tag{10}$$

From time t_i onwards algorithm EA now optimizes with respect to the set A_{i+1}. Note that we only need to know the set A_{i+1} in order to run EA. The set B_{i+1} is only needed in order to check whether we should increase $\tilde{\ell}$. Again applying Chernoff bounds and a union bound over all permissible t we have whp that

$$\text{for all } i \text{ s.t. } t_i \leq T\text{:} \quad d(B_{i+1}, t) = d(A_{i+1}, t) \pm \gamma \quad \text{for all } 1 \leq t \leq \min\{t_{i+1}, T\}. \tag{11}$$

Finally, the probability to flip a linear number of bits in one round is exponentially small, so whp we have

$$\text{for all } 1 \leq t \leq T\text{:} \quad \text{the number of bits flipped in round } t \text{ is less than } \beta\gamma n. \tag{12}$$

This condition implies in particular that $d(A_i, t+1) = d(A_i, t) \pm \gamma$ and $d(B_i, t+1) = d(B_i, t) \pm \gamma$ for all $i \geq 1$, $t \leq T$. Also note that the probability to flip exactly k bits in one round drops at least geometrically

as $k \to \infty$, so the number of zero bits in ξ^t satisfies condition 2 in Theorem 6. Similarly for the number of zero bits in A_i or in B_i at time t. Note that this implies in particular, that we may assume that within the first T steps of algorithm EA the density of any of these sets will never change opposite to the direction of its drift by more than γ.

Assume now that we can show (we will do this below) that algorithm EA with respect to the function $\tilde{f}$ satisfies $d([n],t) > \varepsilon + 3\gamma$ for all $1 \le t \le T$. Then one easily checks that (10) and (11) imply that for all i such that $t_i \le T$ we have $d(B_i,t) > \varepsilon$ for all $t < t_i$ which implies that $\ell(t) = \tilde{\ell}(t)$ for all $t \le T$. Thus it indeed suffices to analyse EA with respect to $\tilde{f}$, and this is what we will now do. More precisely, for a suitably chosen constant $\delta > 0$ we will show by induction that for all i such that $t_i \le T$ the following three properties hold with probability at least $1 - e^{-2\mu n}$:

(i) $d(A_i,t_i) = \varepsilon \pm 2\gamma$,

(ii) $d([n] \setminus (A_i \cup B_{i-1}, t_i) \ge \varepsilon + \delta$,

(iii) $\forall t_i \le t \le \min\{t_{i+1}, T\}$: $d([n],t) \ge \varepsilon + 3\gamma$.

Note that the choice of T together with a union bound argument thus implies that these properties all hold for all i simultaneously, which then concludes the proof.

First consider (i). We know by (11) that $d(A_i,t_i) = d(B_i,t_i) \pm \gamma$. We also argued already that for all $t < t_i$ we have $d(B_i,t) > \varepsilon$. As $d(B_i,t_i) \le \varepsilon$ by definition of t_i, the claim thus follows easily together with (12).

Next we prove (ii) and (iii) together. By induction we may assume that (i) and (ii) hold for $i-1$. (Since (ii) is satisfied for $i = 0$, this will also yield the base case $i = 1$.) By the induction assumption we know that $d([n],t_{i-1}) \ge \alpha(\varepsilon - 2\gamma) + (1-\alpha-\beta)(\varepsilon+\delta) \ge \varepsilon + (1-\alpha)\delta - \beta$ by suitable choice of the constants, cf. (9). By definition of t_{i-1} and (12) we know that $\varepsilon - \gamma \le d(B_{i-1},t_{i-1}) \le \varepsilon$. Let $R = [n] \setminus (A_i \cup B_{i-1})$. By the (random) choice of A_i and (10) we thus know that

$$d(A_i,t_{i-1}) = d([n],t_{i-1}) \pm \gamma \qquad \text{and} \qquad d(R,t_{i-1}) = d([n] \setminus B_{i-1}, t_{i-1}) \pm \gamma \ge d([n],t_{i-1}) \pm \gamma. \quad (13)$$

We first compute the drift within A_{i+1}. To do this observe that the choice of the weights in the function $\tilde{f}$ are such that the bit flips within A_{i+1} dominate: if they lead to a decreased number of zeroes within A_{i+1} then the offspring is accepted regardless of what happens for the bits outside of A_{i+1} (unless the density within B_i drops below $1-\varepsilon$ and the offspring is thus rejected for this reason). Similarly, if the number of zeroes within A_{i+1} increases then the offspring is always rejected. Only in the remaining case the bits outside of A_{i+1} matter. In particular, the density within A_{i+1} is non-increasing.

Consider $t_i' := \inf\{t \in [t_{i-1},T] : d(A_{i+1},t) \le \varepsilon + (1-\alpha)\delta - \beta\}$, and note that $t_i' \le t_i$ by (11). By (12), and since $d(A_{i+1},t_{i-1}) \ge \varepsilon + (1-\alpha)\delta - \beta - \gamma$ by (13), we also have $d(A_{i+1},t_i') \ge \varepsilon + (1-\alpha)\delta - \beta - \gamma$. Since the density in A_{i+1} is non-increasing, for all $t \in [t_i', \min\{t_i,T\}]$ the probability that we flip at least two zero bits in A_{i+1} to a one is bounded by $O(\varepsilon^2)$, implying that the drift within A_{i+1} is essentially determined by the probability that we flip exactly one bit within A_{i+1} (which is $(1+o(1))\alpha c e^{-\alpha c}$) and that this bit is a zero bit (which gives a factor $d(A_{i+1},t)$). Also the probability that we flip a bit within B_{i-1} is bounded by $c\beta$, which we may assume to be also of the order $O(\varepsilon^2)$. We thus know that the expected drift for the density within A_{i+1} is for all $t_i' \le t \le \min\{t_i,T\}$

$$-\frac{(1+o(1))\alpha c e^{-\alpha c} \cdot d(A_{i+1},t) + O(\varepsilon^2)}{\alpha n} = -\frac{\varepsilon c}{n} e^{-\alpha c}(1+O(\varepsilon)) =: \Delta_1.$$

Next we consider what happens within R. By (13) and Lemma 9 we may assume that $d(R,t) \ge d(A_{i+1},t) - 3\gamma$ for all $t_{i-1} \le t \le \min\{t_i,T\}$. In particular, $d(R,t_i') \ge d(A_{i+1},t_i') - 3\gamma \ge \varepsilon + (1-\alpha)\delta - \beta - 4\gamma$. We want to show that the density in $d(R)$ *increases* to at least $\varepsilon + \delta$. Thus, we may assume without loss of generality that $d(R,t) = \varepsilon \pm 2\delta$ for all $t_i' \le t < \min\{t_i,T\}$. We will now compute a lower bound on the drift within R. Consider first the case that the number of zero bits in A_{i+1} remains the same (which happens with probability at most $e^{-\alpha c} + O(\varepsilon)$). In this case we accept the changes within R similarly as we did for those within A_{i+1}, as explained above. In other words: with probability at most $e^{-\alpha c}(|R|/n)ce^{-(|R|/n)c} \cdot \varepsilon + O(\delta)$ the number of zero bits within R decreases by one, and the contribution of events where it may increases is $O(\varepsilon^2)$.

If the number of zero bits within A_{i+1} decreases (which happens with probability $\alpha c e^{-\alpha c}\varepsilon + O(\varepsilon^2)$), then we accept any change within R. In this case the expected change in the number of zero bits within R is $c|R|/n + O(\varepsilon)$. Summarizing, we observe that the expected drift for the density within R is

$$\frac{\alpha c e^{-\alpha c}\varepsilon \cdot \frac{c|R|}{n} - e^{-\alpha c}\frac{|R|}{n} c e^{-(|R|/n)c} \cdot \varepsilon}{|R|} + O(\delta + \varepsilon^2) = \frac{\varepsilon c}{n} e^{-\alpha c}\left(\alpha c - e^{-(1-\alpha)c}\right) + O(\tfrac{1}{n}(\delta + \varepsilon^2 + \beta)) =: \Delta_2,$$

where we used that $|R|/n = 1 - \alpha + O(\beta)$.

Observe that, up to the error terms, the ratio $\Delta_2/|\Delta_1|$ is exactly the left hand side of (8). By making δ, ε and β small enough we may thus assume that

$$\frac{\Delta_2}{|\Delta_1|} > \frac{\alpha}{1-\alpha} + O(\beta + \gamma). \tag{14}$$

We now employ the tail bounds in Theorem 6: as we have $d(A_{i+1}, t_i') \geq \varepsilon + (1-\alpha)\delta - \beta - \gamma$, Theorem 6 tells us that with probability $1 - e^{-\Omega(n)}$ it takes at least $S = [(1-\alpha)\delta - \beta - 3\gamma]/|\Delta_1|$ steps to decrease it to $\varepsilon + \gamma$. Before this density is reached, in each round the density in B_{i+1} is strictly larger than ε by (12). On the other hand, $d(R, t_i') \geq \varepsilon + (1-\alpha)\delta - \beta - 4\gamma$. Again by Theorem 6 the density in R is increased to at least $\varepsilon + \delta + \gamma$ with probability $1 - e^{-\Omega(n)}$ after at most $(\alpha\delta + \beta + 6\gamma)/\Delta_2 \overset{(14)}{<} S$ steps, and stays above $\varepsilon + \delta$ by Theorem 6b. This proves both (ii) and (iii) of the induction and thus concludes the proof of the theorem. □

A natural question is to determine the behaviour of the algorithm if the mutation parameter c lies in the interval $[1, 2.2)$. We leave this as an intriguing open problem.

4.2 Linear Functions

In this section we give a new proof that the optimization time of the EA with mutation parameter c is $(1 + o(1))\frac{e^c}{c} n \log n$ for any constant $c > 0$ if we restrict ourselves to *linear* functions $f(\xi) = \sum_{i=1}^n a_i \xi_i$, as was first proven in [17]. Recall that we assume that the weights are sorted and positive, $a_1 \geq \ldots \geq a_n > 0$.

Theorem 11. *For any linear function f and every constant $c > 0$, algorithm EA with mutation rate c/n has optimization time $(1 + o(1))e^c/c \cdot n \log n$ whp.*

The lower bound in Theorem 11 is rather straightforward, and we only give a sketch. Assume the number X_t of zero bits in ξ^t is small, say $X_t \leq n^{1-\varepsilon}$. Then the probability to flip at least two zero bits in the same round is $O(n^{-2\varepsilon})$. This will turn out to be negligibly small, so let us ignore such "bad" rounds for the moment for the sake of clarity. Then the number of zero bits can only decrease if we flip exactly one zero bit, and no one bits. Let r be the number of rounds in which we flip exactly one position. Then the number of such rounds until no zero bits are left is (if there were no bad rounds) the same as the number of rounds needed in a coupon collector process to reduce the number of missing coupons from from $n^{1-\varepsilon}$ to 0, which is well-known to be at least $(1 - \varepsilon - o(1))n \log n$ whp. On the other hand, in each round we have probability $(1 + o(1))ce^{-c}$ to flip exactly one bit. So the (total) number of rounds until we have seen whp $r = (1 - \varepsilon - o(1))n \log n$ rounds in which we flip only a single bit is $(1 - \varepsilon - o(1))e^c/c \cdot n \log n$, by the Chernoff bounds. It is easy to adapt the argument to also cope with bad rounds, e.g., by dividing the process into phases form $n^{1-\varepsilon}$ to $n^{1-2\varepsilon}$, from $n^{1-2\varepsilon}$ to $n^{1-4\varepsilon}$ etc., and observing that in each phase with very high probability the number of bad rounds is so small that it does not significantly change the time to reach the next phase.

For the upper bound, we split the proofs in two steps: we need a linear time until there are less than εn zero bits left (Lemma 13), which are then removed in an additional $(1 + o(1))e^c/c \cdot n \log n$ steps (Lemma 12).

Lemma 12. *Let $c > 0$ be a constant. Then there is a constant $\varepsilon > 0$ such that the following holds for every $\xi_0 \in \{0,1\}^n$ that has at most εn zero bits. If algorithm EA with mutation rate c/n starts with initial string ξ^0 then whp it finds the optimum of any n-bit linear function with positive weights after at most $(1 + o(1))e^c/c \cdot n \log n$.*

Proof. Our aim is to apply Theorem 5. For the function g we count the number of zeroes, but we also distinguish between the positions where the zeroes occurs, giving those with a higher coefficient a_i also a higher weight. It seems natural to use the a_i themselves as weights. However, this does not work since Theorem 3 only gives a very bad bound if, for example, $a_1/a_n \gg n$. Instead, we fix a constant $C > 0$, to be chosen later, and we use the objective function

$$g(\xi) := 2^C \cdot \sum_{1 \le i \le n, \xi_i = 0} 2^{-Ci/n}.$$

Let X_t be the number of zero bits in ξ^t, and observe that $X_t \le g(\xi^t) \le 2^C X_t$.

We want to compute the drift $\mathbb{E}[g(\xi^{t+1}) \mid \xi^t = \xi] - g(\xi)$, so assume $\xi^t = \xi$. First observe that the probability to flip at least two zero bits in the same round is $O((X_t/n)^2) = O(g(\xi)^2/n^2)$, so this case contributes at most $O(g(\xi)^2/n^2)$ to the drift (which, as we will see, is negligible). On the other hand, if no zero bit is flipped then the offspring is rejected. So it remains to consider the case where exactly one zero bit is flipped (and possibly additional one bits are flipped). We call the index of this bit k. If moreover no one bits are flipped, then this happens with probability $(1 + o(1))ce^{-c}/n$, and we accept the offspring. Hence, this case contributes

$$-(1 + o(1))\frac{ce^{-c}}{n} 2^C \cdot \sum_{1 \le k \le n, \xi_k = 0} 2^{-Ck/n} = -(1 + o(1))\frac{ce^{-c}}{n} g(\xi)$$

to the drift. It remains to consider the case where also one bits are flipped. If all one bits that are flipped have an index i that satisfies $i > k$ then the offspring may or may not be accepted, but in any case the function g increases by at most $S_k := 2^C \sum_{j>k, \xi_j \text{ flipped}} 2^{-Cj/n}$. On the other hand hand, if a one bit i is flipped with $i < k$ then the fitness does not increase, so the offspring is rejected. Hence, the change in g is 0. Note that S_k is always non-negative, so we can upper bound the positive contribution to the drift from all cases where we flip exactly one zero bit by

$$\sum_{1 \le k \le n, \xi_k = 0} \Pr[\xi_k \text{ flips}] \cdot \mathbb{E}[S_k] = \sum_{1 \le k \le n, \xi_k = 0} \frac{c}{n} \cdot 2^C \sum_{i=k+1}^{n} \frac{c}{n} 2^{-Ci/n} \le \frac{c2^C}{n} \sum_{1 \le k \le n, \xi_k = 0} 2^{-Ck/n} \underbrace{\sum_{i=1}^{\infty} \frac{c}{n} 2^{-Ci/n}}_{= c/(C \log 2) + o(1)}$$

$$= (1 + o(1))\frac{c^2}{nC \log 2} g(\xi).$$

Assume that C is so large that $\delta := ce^c/(C \log 2) \le 1/3$. Then, bringing all cases together, we obtain for any string ξ,

$$\mathbb{E}[g(\xi^{t+1}) \mid \xi^t = \xi] \le g(\xi) \left(1 - (1 + o(1))\frac{(1 - \delta)ce^{-c}}{n} + O\left(\frac{g(\xi)}{n^2}\right)\right), \tag{15}$$

where the term $O(g(\xi)/n^2)$ covers the case that at least two zero bits are flipped. This inequality implies two things. Firstly, we apply (15) for $C = 3ce^c/\log 2$ (i.e., $\delta = 1/3$). Since $g(\xi^0) \le 2^C X_0 \le 2^C \varepsilon n$, we may choose ε such that the error term $O(g(\xi^0)/n^2)$ is at most $\frac{1}{3}ce^{-c}/n$ at time $t = 0$. Then by Theorem 6, for every $\gamma > 0$ whp $X_t = O(g(\xi^t)/n)$ becomes smaller than γ in time $t_0 = O(n)$, and it stays so for an exponential number of steps. Having proven this, we may now apply (15) for any constant $C > 3ce^c/\ln 2$ (and thus, any constant $0 < \delta < 1/3$). Since then $g(\xi^t) \le 2^C X_t = O(2^C \gamma)$, we may choose γ small enough so that the error term $O(g(\xi^t)/n^2)$ is at most $\delta ce^{-c}/n$ for $t \ge t_0$. Thus we obtain $\mathbb{E}[g(\xi^{t+1}) \mid \xi^t = \xi] \le g(\xi)(1 - (1 - 3\delta)ce^{-c}/n)$. Then by Theorem 5, whp we reach zero after at most $(1 + o(1))\frac{ne^c}{(1-3\delta)c} \log(g(\xi^{t_0}))$ additional steps. Since $g(\xi^{t_0}) = O(n)$, and since we can choose δ arbitrarily small, this proves that whp the algorithm finds the optimum after $(1+o(1))e^c/c \cdot n \log n$ steps. □

It remains to show that the number of zeroes is at most εn after a short (i.e., linear) time.

Lemma 13. *For any $c, \varepsilon > 0$ there is $C > 0$ such that after Cn steps of algorithm EA on any n-bit linear function with positive weights, whp there are at most εn zero bits left.*

Even though Lemma 13 implies that in absolute terms the first phase of an EA takes only linear time, while the second phase requires an additional log-factor, the proof of Lemma 13 turns out particularly tricky. We will use an induction-type argument, but not over natural numbers, but rather over values of c. For technical reasons, we fix $c_{\max} > 0$ sufficiently large (but otherwise arbitrary) and prove the statement for all $0 < c < c_{\max}$. Note that since $c_{\max}$ is arbitrary, this nevertheless implies the desired result for all constants $c > 0$.

For the inductive argument to go through, we need to use a slightly stronger statement, which we spell out in Lemma 14. We will allow two types of noise. We say that there is δ_1-*noise of type 1*, if in each round we flip a coin, and with probability δ_1 an adversary may decide whether we go to the offspring or not, regardless of the objective values. We say that there is δ_2-*noise of type 2* if in each round we flip a coin, and with probability $1 - \delta_2$ an adversary may decide on an arbitrary non-negative (!) penalty by which the fitness of the offspring is reduced. Note that the adversary has to set the penalty without looking at the outcome of the offspring.

Note that Lemma 14 contains Lemma 13 as a special case (the case without noise), so proving Lemma 14 also concludes the proof of Theorem 11.

Lemma 14. *If $c_{\max} > 0$ is a sufficiently large constant, then for all $0 < \varepsilon \leq 1$ and all $0 < c < c_{\max}$ there is $C > 0$ such that for all $n \geq 1$ the following holds. Consider algorithm EA with arbitrary starting string, with mutation rate c/n, and with δ_1-noise of type 1 and δ_2-noise of type 2, where $\delta_1 = \delta_1(\varepsilon, c) := \varepsilon \exp\{-ce^{4c_{\max}+e^{5c_{\max}}}\}$ and $\delta_2 = \delta_2(c) := e^{2(c-c_{\max})}$. Then for any n-bit linear function with positive weights, whp for all $t \in [Cn, e^{n/C}]$ the t-th search point ξ^t has at most εn zero bits.*

Proof. We first investigate the case that c is small. Observe that for a fixed $0 < \varepsilon \leq 1$, and c sufficiently small the probability to flip more than one bit in the same round can be bounded by $O(c^2)$. Thus, if the number of zero bits is at least $\varepsilon n/2$ then the drift is at most $-\delta_2 ce^{-c}\varepsilon/4 + \delta_1 c + O(c^2)$. Here $ce^{-c}\varepsilon/2$ is a lower bound (for large n) for the probability that we flip exactly one bit and that this bit is a zero bit and the factor $\delta_2 - \delta_1 \geq \delta_2/2$ is a lower bound for the probability that the noise of type 1 or 2 will not obstruct us from accepting this bit flip. The remaining two terms are upper bounds for the cases in which the number of zero bits increases. That is, if $c_{\max}$ is sufficiently large then for $c = \varepsilon e^{-4c_{\max}}$ this drift is negative, and the lemma follows for the pair ε, c from Theorem 6. In other words, for each fixed $0 < \varepsilon \leq 1$ the lemma holds for some $c > 0$.

Assume now for the sake of contradiction that the lemma is false for some $0 < \varepsilon \leq 1$, and let $\varepsilon_0 > 0$ be the supremum over all ε for which it is false. Then there exists an $\varepsilon > \varepsilon_0/2$ for which the lemma is false. From now on we will stick with this ε and derive a contradiction. Let $c_0 = c_0(\varepsilon)$ be the infimum over all c for which the lemma fails for the pair (ε, c). Then $c_0 \geq \varepsilon e^{-4c_{\max}}$ by the considerations above. We will show that there exists $\rho_c = \rho_c(\varepsilon) > 0$ such that the lemma holds for all pairs (ε, c) with $c_0 \leq c < \max\{c_0(1 + \rho_c), c_{\max}\}$, thus yielding the desired contradiction. In the following we assume that c is an arbitrary but fixed value in the interval $[c_0, \max\{c_0(1 + \rho_c), c_{\max}\})$, where ρ_c is defined after equation (18) below.

Before going into details we explain the main idea of our proof. First note that the theorem is *true* for $\varepsilon^+ := 2\varepsilon$: if $2\varepsilon > 1$ this is trivial, otherwise this follows from the definition of ε_0 and the fact that $\varepsilon \geq \varepsilon_0/2$. Thus we know that algorithm EA (with parameter $p = c/n$) reduces the number of zeroes in linear time to $2\varepsilon n$ (and will stay below this number for an exponential time). We will show that there exists $\alpha > 0$ and $0 < \hat{\varepsilon} < \varepsilon$ so that after an additional linear number of steps, there are at most $\hat{\varepsilon}\alpha n$ zero bits left among the first αn bits. The remaining part has only $\tilde{n} := (1-\alpha)n$ bits, thus the mutation rate in this part is $(1-\alpha)c/\tilde{n}$. We will choose α in such a way that $\tilde{c} := (1-\alpha)c < c_0$, so we know from the choice of c_0 that the lemma is true for the pair $(\varepsilon, \tilde{c})$. Thus, the remaining string contains at most $\varepsilon\tilde{n}$ zeros after linear time, at least if we would apply the algorithm only to that part of the string. In order to really derive a contradiction we must take into account that the αn part of the string provides noise for the optimization of the remaining part: if a zero bit in the first part is flipped into a one bit (and possibly other bits in this part are also flipped), then this gives noise of type 1; if one or several one bits are flipped then this may give noise of type 2. We thus need to check that both types of noise are sufficiently small. More precisely, we will check that this additional noise is small enough, so that the total noise is less than the values for δ_1 and δ_2 with respect to ε and $\tilde{c}$.

To prove these claims we need to define some constants. Let $\delta_1 := \delta_1(\varepsilon, c)$ and $\delta_2 := \delta_2(c)$ be as

defined in the statement of the lemma. Moreover, let

$$\hat{\varepsilon} := e^{4c_{\max}+e^{5c_{\max}}}\delta_1 \qquad \text{and} \qquad \beta := e^{-4c_{\max}}\hat{\varepsilon}. \tag{16}$$

Recall that $\varepsilon e^{-4c_{\max}} \leq c_0 \leq c < c_{\max}$. So we may assume without loss of generality that $c_{\max}$ is large enough so that

$$(2e\varepsilon c \exp\{-e^{5c_{\max}}\})^{\exp\{e^{5c_{\max}}\}} \leq \delta_1 \quad \text{and} \quad \max\{\delta_1, 4c\beta\} \leq \tfrac{1}{20}\hat{\varepsilon}\delta_2 e^{-c_{\max}} \leq \tfrac{1}{4} \quad \text{and} \quad \beta \leq \tfrac{1}{4}. \tag{17}$$

Note that (17) implies in particular that

$$1 - c\beta \geq e^{-2c\beta} \qquad \text{and} \qquad 1 - 2c\beta \geq e^{-4c\beta}. \tag{18}$$

With these definitions at hand, we let ρ_c be so small that $c(1-\beta) < c_0$ whenever $c_0 \leq c \leq c_0(1+\rho_c)$. Note that β implicitly depends on c here, but it is easy to check that the definition $\rho_c := \min_{\varepsilon e^{-4c_{\max}} \leq c \leq c_{\max}}\{\beta(c)\}$ yields the desired inequalities.

As before, we may assume that the bits are sorted by weights, in descending order. We call the first βn bits *the first region* R_1, the next βn bits *the second region* R_2, and the remaining $(1-2\beta)n$ bits the *third region* R_3. We assume further without loss of generality that the maximum weight in the second section is 1 (by scaling all n weights by the same factor), and we split the second and third region into blocks as follows. For $i \geq 1$, the block B_i consists of all bits in $R_2 \cup R_3$ with weight in $[(1+\mu)^{-i-1}, (1+\mu)^{-i})$. Note that some of the blocks may be empty. Recall that $d(B_i) = d(B_i, t)$ is the density of zero bits in B_i with respect to ξ^t, for any non-empty block B_i, and similarly for $d(R_1)$ etc.

We distinguish two cases (in which we will choose $\alpha = \beta$ and $\alpha = 2\beta$, respectively) depending on how many blocks there are in the second region. Note that we also count empty blocks here, so the number of blocks in the second region is $\lceil \log_{1+\mu} 1/a_i \rceil$, where $i = 2\beta n$ is the last index in the second region. To describe the cases we need two more constants. Let

$$D := \min\{D \in \mathbb{N} \mid (1+\mu)^D \geq \exp\{e^{5c_{\max}}\}\}, \quad \text{where } \mu := e^{-c_{\max}} \exp\{-2e^{5c_{\max}}\}\delta_2\beta\hat{\varepsilon}/32. \tag{19}$$

First case: there are more than D blocks in R_2. In this case we will set $\alpha = \beta$. Recall from above that we may assume that the number of zero bits is at most $2\varepsilon n$. We consider the effect of the algorithm on R_1, i.e. on the first βn bits. We want to argue that we can bring $d(R_1)$ below $\hat{\varepsilon}$ in linear time, and that it stays so for exponential time. To this end, we will show that we have a negative drift whenever $d(R_1) \geq \hat{\varepsilon}/2$, so let us assume the latter. The probability that we flip exactly one zero bit in R_1, and no other bit anywhere, is at least $ce^{-c_{\max}}\beta\hat{\varepsilon}/2$ if n is sufficiently large; in this case we accept this bit flip with probability at least $\delta_2 - \delta_1 > \delta_2/2$, and the number of zero bits decreases by one. On the other hand, we can increase the number of zero bits in R_1 only if one of the following (non-disjoint) events occurs:

(i) at least two bits in R_1 are flipped: this happens with probability at most $c^2\beta^2$ and the expected increase is then at most $\beta c + 2 \leq 3$, with room to spare;

(ii) at least one bit is flipped in R_2: this happens with probability at most $c\beta$ and the expected increase (in R_1) is then at most βc;

(iii) more than $(1+\mu)^D$ zero bits are flipped in R_3: since there are at most $2\varepsilon n$ zero bits, this probability is bounded by the probability that a binomially distributed random variable with parameters $2\varepsilon n$ and c/n is at least $(1+\mu)^D \geq \exp\{e^{5c_{\max}}\}$; using the Chernoff bound $\Pr[\mathrm{Bin}(N,q) \geq (1+\delta)Nq] \leq (e/(1+\delta))^{(1+\delta)Nq}$ with $Nq = 2\varepsilon c$ and $(1+\delta)Nq = \exp\{e^{5c_{\max}}\}$ we see that this probability is at most $(2e\varepsilon c \cdot \exp\{-e^{5c_{\max}}\})^{\exp\{e^{5c_{\max}}\}}$; the expected increase in this case is again at most βc;

(iv) there is noise of type 1: this happens with probability δ_1 and also results in an expected increase of at most βc.

Summarizing the above terms, we see that the expected drift for the number of zero bits in the first region can be bounded by

$$\mathbb{E}[X_{t+1} - X_t] \leq -\tfrac{1}{4}\delta_2 c e^{-c_{\max}}\beta\hat{\varepsilon} + 3c^2\beta^2 + c^2\beta^2 + (2e\varepsilon c \exp\{-e^{5c_{\max}}\})^{\exp\{e^{5c_{\max}}\}}\beta c + c\beta\delta_1 \overset{(17)}{=} -\Omega(1).$$

Hence, by Theorem 6 after linear time $d(R_1) \leq \hat{\varepsilon}/2$, and since we have negative drift whenever $d(R_1)$ is in the interval $[\hat{\varepsilon}/2, \hat{\varepsilon}]$, it will stay below $\hat{\varepsilon}$ for exponential time.

We want to apply the lemma for ε and $\tilde{c} = c(1-\beta) < c_0$ to conclude that in this period $d(R_2 \cup R_3)$ drops below ε. So we must show that the additional noise caused by the first region is not too large. Noise of type 1 is generated whenever a zero bit in R_1 is flipped, which happens with probability at most $c\beta\hat{\varepsilon}$. Thus the total amount of noise of type 1 is at most

$$\delta_1(\varepsilon, c) + c\beta\hat{\varepsilon} \overset{(16)}{\leq} \delta_1(\varepsilon, c) + c\beta \frac{-\partial\delta_1(\varepsilon, c)}{\partial c} \leq \delta_1(\varepsilon, \tilde{c}),$$

where in the last step we used convexity of δ_1. For noise of the type 2, the fitness of the offspring may decrease if any bit in R_1 is flipped, which happens with probability at most $c\beta$. Hence, the overall probability that the adversary will *not* add a penalty is at least

$$\delta_2(c) \cdot (1 - c\beta) \overset{(18)}{\geq} \delta_2(\tilde{c}) e^{2c\beta} \cdot e^{-2c\beta} \geq \delta_2(\tilde{c}),$$

as required. This concludes the first case.

Second case: there are at most D blocks in the second region. In this case we will apply the lemma to the third region ($\alpha = 2\beta$) with parameters ε and $\tilde{c} = c(1-2\beta) < c_0$. To bound the noise coming from R_1 and R_2, we will first show that whp $d(R_1 \cup R_2) \leq \hat{\varepsilon}$ after linear time. We use the auxiliary fitness function

$$\phi(\xi) := \sum_{i \geq 1} |\{j \in B_i : \xi_j = 0\}| \cdot (1+\mu)^{-i},$$

and we let $\phi_{\max}$ be the fitness of the zero string, i.e. $\phi_{\max} = \sum_{i \geq 1} |B_i| \cdot (1+\mu)^{-i}$. Note that by definition of the second case $\phi_{\max} \geq \beta \exp\{e^{-5c_{\max}}\}$. We claim that whp whenever $d(R_2, t) > \hat{\varepsilon}/2$ then $\phi(\xi^t) > \hat{\varepsilon}(1+\mu)^{-D}\phi_{\max}/4$, from which a negative drift for $\phi(\xi^t)$ will result. Hence, by Theorem 6 it follows that after linear time $d(R_2) \leq \frac{3}{4}\hat{\varepsilon}$, and that this bounds remains true for an exponential number of steps. By Lemma 9 the density in R_1 is then also at most $\hat{\varepsilon}$, so in total $d(R_1 \cup R_2) \leq \hat{\varepsilon}$.

Let us first show the claimed lower bound on $\phi(\xi^t)$, so assume $d(R_2, t) > \hat{\varepsilon}/2$. We call a block in the third region *long* if it has length at least $\beta\mu n/2$, and *short* otherwise. The contribution of short blocks to $\phi_{\max}$ is at most

$$\sum_{B_i \text{ short}} |B_i|(1+\mu)^{-i} < \frac{1}{2}\beta\mu n \sum_{i \geq 0} (1+\mu)^{-i} < \beta n. \tag{20}$$

On the other hand, since long blocks have linear length, there can only be a constant number of long blocks. Thus we may apply Lemma 9 to conclude that (after a linear number of steps) we have $d(B_i, t) \geq d(R_2, t)/2$ for all long blocks B_i and for exponential time. Note that the contribution of a block B_i to $\phi(\xi^t)$ can be written as $d(B_i, t)|B_i|(1+\mu)^{-i}$, and the contribution of R_2 is at least $d(R_2, t)\beta n(1+\mu)^{-D}$, where the later follows from the definition of the second case. Hence, if $d(R_2, t) \geq \hat{\varepsilon}/2$ then

$$\phi(\xi^t) \geq \frac{d(R_2, t)\beta n}{(1+\mu)^D} + \sum_{B_i \text{ long}} \frac{1}{2} d(R_2, t)|B_i|(1+\mu)^{-i} > \frac{\hat{\varepsilon}}{4(1+\mu)^D} \left(2\beta n + \sum_{B_i \text{ long}} |B_i|(1+\mu)^{-i} \right)$$
$$\overset{(20)}{\geq} \frac{\hat{\varepsilon}\phi_{\max}}{4(1+\mu)^D},$$

proving the first claim from above.

It remains to show that ϕ has a negative drift for $\phi(\xi^t) > \hat{\varepsilon}(1+\mu)^{-D}\phi_{\max}/4$. First observe that with probability $ce^{-c}(1+o(1))$ there is exactly one bit flipped, and there is no noise with probability at least $\delta_2 - \delta_1 > \delta_2/2$. In this case each bit has probability $1/n$ to be the flipped one, so in this case we decrease $\phi(\xi^t)$ by $\phi(\xi^t)/n$ on average. Thus the contribution to the drift is (for sufficiently large n) at least $-ce^{-c}\delta_2\phi(\xi^t)/(4n)$. If at least one bit in R_1 is flipped (which happens with probability at most $c\beta$), then this increases $\phi(\xi^t)$ in expectation by at most $c\phi_{\max}/n$, so this adds at most $c^2\beta\phi_{\max}/n$. Similarly, noise of type 1 contributes at most $c\delta_1\phi_{\max}/n$ to the drift. Finally, consider the case where no bit in R_1 flips, and where no noise of type 1 occurs. Let r_i and s_i be the number of bits in the i-th block B_i that flip from zero to one and from one to zero, respectively. Then a necessary (but not sufficient) condition for accepting the offspring is $\sum_i s_i(1+\mu)^{-i-1} - \sum r_i(1+\mu)^{-i} \leq 0$. Note that this condition is still necessary if noise of type 2 occurs. Thus, in this case $\Delta\phi = \sum_i s_i(1+\mu)^{-i} - \sum r_i(1+\mu)^{-i} \leq \mu \sum_i r_i(1+\mu)^{-i} \leq \mu \sum_i r_i$, which is at

most $c\mu$ in expectation. Summarizing, we see that for $d(R_2,t) \geq \hat{\varepsilon}/2$ (and thus $\phi(\xi^t) \geq \hat{\varepsilon}(1+\mu)^{-D}\phi_{\max}/4$) the drift is at most

$$\mathbb{E}[\phi(\xi^{t+1}) - \phi(\xi^t)] \leq -ce^{-c}\delta_2\hat{\varepsilon}(1+\mu)^{-D}\phi_{\max}/(16n) + c^2\beta\phi_{\max}/n + c\mu + c\delta_1\phi_{\max}/n \overset{(17),(19)}{=} -\Omega(1)$$

for sufficiently large $c_{\max}$.

Therefore $d(R_1 \cup R_2) \leq \hat{\varepsilon}$ for an exponential time. We now apply the lemma similarly as in the first case, this time to R_3. I.e., we set $\alpha = 2\beta$, and apply the lemma for ε and $\tilde{c} = c(1-2\beta) < c_0$. We need to check that we do not introduce too much noise. For noise of type 1, we have at most

$$\delta_1(\varepsilon,c) + 2c\beta\hat{\varepsilon} \overset{(16)}{\leq} \delta_1(\varepsilon,c) + 2c\beta\frac{-\partial\delta_1(\varepsilon,c)}{\partial c} \leq \delta_1(\varepsilon,\tilde{c}),$$

where again we used convexity in the last step. For noise of the type 2 we conclude as in the first case that the probability *not* to have noise is at least

$$\delta_2(c)(1-2c\beta) \overset{(18)}{\geq} \delta_2(\tilde{c})e^{4c\beta}\cdot e^{-4c\beta} \geq \delta_2(\tilde{c}),$$

as desired. □

Acknowledgement. After being appointed at ETH Zürich, Jiří Matoušek convinced his colleagues in the algorithms group to establish a joint seminar to teach undergraduate students to do (and love) research. In the third edition of this seminar the second author of this paper presented the topic at hand. More precisely, the task was to come up with fresh ideas to elegantly reprove the state of the art. While this turned out to be a too hard task for undergraduate students, it eventually led to the present paper, which, at least for us, will always be a memory for our much admired colleague Jirka Matoušek.

References

[1] Thomas Bäck. *Evolutionary algorithms in theory and practice: evolution strategies, evolutionary programming, genetic algorithms*. Oxford University Press, 1996.

[2] Benjamin Doerr and Leslie Goldberg. Adaptive drift analysis. *Algorithmica*, pages 1–27, 2011.

[3] Benjamin Doerr, Thomas Jansen, Dirk Sudholt, Carola Winzen, and Christine Zarges. Mutation rate matters even when optimizing monotonic functions. *Evolutionary computation*, 21(1):1–27, 2013.

[4] Benjamin Doerr, Daniel Johannsen, and Carola Winzen. Multiplicative drift analysis. *Algorithmica*, 64(4):673–697, 2012.

[5] Benjamin Doerr, Daniel Johannsen, and Carola Winzen. Non-existence of linear universal drift functions. *Theoretical Computer Science*, 436:71–86, 2012.

[6] Stefan Droste, Thomas Jansen, and Ingo Wegener. On the analysis of the (1+1) evolutionary algorithm. *Theoretical Computer Science*, 276:51–81, 2002.

[7] Bruce Hajek. Hitting-time and occupation-time bounds implied by drift analysis with applications. *Advances in Applied Probability*, 13:502–525, 1982.

[8] Jun He and Xin Yao. Drift analysis and average time complexity of evolutionary algorithms. *Artificial Intelligence*, 127:57–85, 2001.

[9] Jun He and Xin Yao. A study of drift analysis for estimating computation time of evolutionary algorithms. *Natural Computing*, 3:21–35, 2004.

[10] Jens Jägersküpper. Combining Markov-chain analysis and drift analysis. *Algorithmica*, 59(3):409–424, 2011.

[11] Thomas Jansen. On the brittleness of evolutionary algorithms. In *Foundations of Genetic Algorithms*, pages 54–69. Springer, 2007.

[12] Daniel Johannsen. *Random Combinatorial Structures and Randomized Search Heuristics*. PhD thesis, Universität des Saarlandes, 2010.

[13] Timo Kötzing. Concentration of first hitting times under additive drift. In *Proceedings of the 2014 conference on Genetic and evolutionary computation*, pages 1391–1398. ACM, 2014.

[14] Boris Mitavskiy, Jonathan Rowe, and Chris Cannings. Theoretical analysis of local search strategies to optimize network communication subject to preserving the total number of links. *International Journal of Intelligent Computing and Cybernetics*, pages 243–284, 2009.

[15] H. Mühlenbein. How genetic algorithms really work: Mutation and Hillclimbing. In R. Männer and R. Manderick, editors, *Parallel Problem Solving from Nature PPSN II*, pages 15–25. North-Holland, Amsterdam, 1992.

[16] Pietro Simone Oliveto and Carsten Witt. Simplified drift analysis for proving lower bounds inevolutionary computation. *Algorithmica*, 59(3):369–386, 2011.

[17] Carsten Witt. Tight bounds on the optimization time of a randomized search heuristic on linear functions. *Combinatorics, Probability & Computing*, 22(2):294–318, 2013.